Refine Search

Search Results -

Terms	Documents
L6 and cach$4	200

Database:
US Pre-Grant Publication Full-Text Database
US Patents Full-Text Database
US OCR Full-Text Database
EPO Abstracts Database
JPO Abstracts Database
Derwent World Patents Index
IBM Technical Disclosure Bulletins

Search: L7

Refine Search | Recall Text | Clear | Interrupt

Search History

DATE: Saturday, May 22, 2004 Printable Copy Create Case

Set Name side by side	Query	Hit Count	Set Name result set
DB=USPT; PLUR=YES; OP=ADJ			
L7	L6 and cach$4	200	L7
L6	L5 and l2	512	L6
L5	hit near6 way	1441	L5
L4	L3 and (way near8 (compar$4 or select$4))	2	L4
L3	L2 and l1	3	L3
L2	way near8 (position$4 or order$4 or arrang$6 or sequen$4)	237174	L2
L1	(4334289 or 4783735 or 5140690 or 5325511 or 5845320 or 5717916 or 6098152 or 6205519).pn.	8	L1

END OF SEARCH HISTORY

h e b b cg b e e ch

Your search matched **0** of **1040503** documents.
A maximum of **500** results are displayed, **15** to a page, sorted by **Relevance Descending** order.

Refine This Search:
You may refine your search by editing the current search expression or enter new one in the text box.

((cache or way) <near/4> replac$) and ((lock$ or unlock Search

☐ Check to search within this result set

Results Key:
JNL = Journal or Magazine **CNF** = Conference **STD** = Standard

---

Results:
No documents matched your query.

h eee e eee g e ch e ch e e e c che

IEEE Xplore® RELEASE 1.7

Welcome
United States Patent and Trademark Office

Try our New Full-text Search Prototype GO

Help

1) Enter a single keyword, phrase, or Boolean expression.
Example: acoustic imaging (means the phrase acoustic imaging plus any stem variations)
2) Limit your search by using search operators and field codes, if desired.
Example: optical <and> (fiber <or> fibre) <in> ti
3) Limit the results by selecting Search Options.
4) Click Search. See Search Examples

```
((cache or way) <near/4>
replac$) and ((lock$ or
unlock$) <near/8> ((first or
last) <near/4> replac$))
```

Start Search Clear

Note: This function returns plural and suffixed forms of the keyword(s).

Search operators: <and> <or> <not> <in> More

Field codes: au (author), ti (title), ab (abstract), jn (publication name), de (index term) More

Search Options:

Select publication types:

- [x] IEEE Journals
- [x] IEE Journals
- [x] IEEE Conference proceedings
- [x] IEE Conference proceedings
- [x] IEEE Standards

Select years to search:

From year: All to Present

Organize search results by:

Sort by: Relevance
In: Descending order

List 15 Results per page

h eee e eee g e ch e ch

Terms used
cache or **way** **near/4** **replac** and **lock** or **unlock** **near/8** **first** or **last** **near/4** **replac**

Found **6,199** of **134,837**

Sort results by relevance

Display results expanded form

Save results to a Binder

Search Tips

Open results in a new window

Try an Advanced Search

Try this search in The ACM Guide

Results 1 - 20 of 200 Result page: **1** 2 3 4 5 6 7 8 9 10 next

Best 200 shown

Relevance scale

1 Architectures: A perspective on the future of massively parallel computing: fine-grain vs. coarse-grain parallel models comparison & contrast

Predrag T. Tosic

April 2004 **Proceedings of the first conference on computing frontiers on Computing frontiers**

Full text available: pdf(277.49 KB) Additional Information: full citation, abstract, references, index terms

Models, architectures and languages for *parallel computation* have been of utmost research interest in computer science and engineering for several decades. A great variety of parallel computation models has been proposed and studied, and different parallel and distributed architectures designed as some possible ways of harnessing parallelism and improving performance of the general purpose computers.*Massively parallel connectionist models* such as *artificial neural networks* (...

Keywords: cellular automata, distributed systems, massively parallel computing, multiprocessor computers, neural networks, parallel computation models

2 Calculating the Eigenvectors of Diagonally Dominant Matrices

M. M. Blevins, G. W. Stewart

April 1974 **Journal of the ACM (JACM)**, Volume 21 Issue 2

Full text available: pdf(655.03 KB) Additional Information: full citation, abstract, references, index terms

An algorithm is proposed for calculating the eigenvectors of a diagonally dominant matrix all of whose elements are known to high relative accuracy. Eigenvectors corresponding to pathologically close eigenvalues are treated by computing the invariant subspace that they span. If the off-diagonal elements of the matrix are sufficiently small, the method is superior to standard techniques, and indeed it may produce a complete set of eigenvectors with an amount of work proportional to the squar ...

3 Special system-oriented section: the best of SIGMOD '94: Sleepers and workaholics: caching strategies in mobile environments (extended version)

Daniel Barbará, Tomasz Imieliński

October 1995 **The VLDB Journal — The International Journal on Very Large Data Bases**, Volume 4 Issue 4

Full text available: pdf(1.73 MB) Additional Information: full citation, abstract, references, citings

h c g e cf c

In the mobile wireless computing environment of the future, a large number of users, equipped with low-powered palmtop machines, will query databases over wireless communication channels. Palmtop-based units will often be disconnected for prolonged periods of time, due to battery power saving measures; palmtops also will frequently relocate between different cells, and will connect to different data servers at different times. Caching of frequently accessed data items will be an important techni ...

Keywords: caching, data management, information services, wireless

4 Cache decay: exploiting generational behavior to reduce cache leakage power

Stefanos Kaxiras, Zhigang Hu, Margaret Martonosi

May 2001 **ACM SIGARCH Computer Architecture News , Proceedings of the 28th annual international symposium on Computer architecture**, Volume 29 Issue 2

Full text available: pdf(1.17 MB) Publisher Site

Additional Information: full citation, abstract, references, citings, index terms

Power dissipation is increasingly important in CPUs ranging from those intended for mobile use, all the way up to high-performance processors for high-end servers. While the bulk of the power dissipated is dynamic switching power, leakage power is also beginning to be a concern. Chipmakers expect that in future chip generations, leakage's proportion of total chip power will increase significantly.

This paper examines methods for reducing leakage power within the cache memori ...

5 Sort-last parallel rendering: Parallel rendering with k-way replication

Rudrajit Samanta, Thomas Funkhouser, Kai Li

October 2001 **Proceedings of the IEEE 2001 symposium on parallel and large-data visualization and graphics**

Full text available: pdf(587.04 KB)

Additional Information: full citation, abstract, references, citings, index terms

With the recent advances in commodity graphics hardware performance, PC clusters have become an attractive alternative to traditional high-end graphics workstations. The main challenge is to develop parallel rendering algorithms that work well within the memory constraints and communication limitations of a networked cluster. Previous systems have required the entire 3D scene to be replicated in memory on every PC. While this approach can take advantage of view-dependent load balancing algorithm ...

Keywords: Parallel rendering, cluster computing, computer graphics systems, interactive visualization

6 On-line restricted caching

Mark Brehob, Richard Enbody, Eric Torng, Stephen Wagner

January 2001 **Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms**

Full text available: pdf(718.78 KB)

Additional Information: full citation, abstract, references, citings, index terms

We study the on-line caching problem in a *restricted* cache where each memory item can be placed in only a restricted subset of cache locations. Examples of restricted caches in practice include victim caches, assist caches, and skew caches. To the best of our knowledge, all previous on-line caching studies have considered on-line caching in *identical* or *fully-associative* caches where every memory item can be placed in any cache location.

In this paper, we focus on ...

h c g e cf c

7 Integrating pointer variables into one-way constraint models
Brad Vander Zanden, Brad A. Myers, Dario A. Giuse, Pedro Szekely
June 1994 **ACM Transactions on Computer-Human Interaction (TOCHI)**, Volume 1 Issue 2

Full text available: pdf(3.71 MB) Additional Information: full citation, abstract, references, citings, index terms

Pointer variables have long been considered useful for constructing and manipulating data structures in traditional programming languages. This article discusses how pointer variables can be integrated into one-way constraint models and indicates how these constraints can be usefully employed in user interfaces. Pointer variables allow constraints to model a wide array of dynamic application behavior, simplify the implementation of structured objects and demonstrational systems, and improve ...

Keywords: Garnet, constraints, development tools, incremental algorithms

8 Memory hierarchies: Direct load: dependence-linked dataflow resolution of load address and cache coordinate
Byung-Kwon Chung, Jinsuo Zhang, Jih-Kwon Peir, Shih-Chang Lai, Konrad Lai
December 2001 **Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture**

Full text available: pdf(1.38 MB) Additional Information: full citation, abstract, references, citings
Publisher Site

An increasing cache latency in future processors incurs profound performance impacts in spite of advanced out-of-order execution techniques. In this paper, we describe an early address resolution mechanism that accurately resolves both regular and irregular load addresses. The basic idea is to build dynamic dependence links from the instruction that updates the base register to the consumer load instructions. Once a new base address is available, it triggers calculations of the new load addresse ...

9 Design of a high-performance ATM firewall
Jun Xu, Mukesh Singhal
August 1999 **ACM Transactions on Information and System Security (TISSEC)**, Volume 2 Issue 3

Full text available: pdf(143.19 KB) Additional Information: full citation, abstract, references, index terms

A router-based packet-filtering firewall is an effective way of protecting an enterprise network from unauthorized access. However, it will not work efficiently in an ATM network because it requires the termination of end-to-end ATM connections at a packet-filtering router, which incurs huge overhead of SAR (Segmentation and Reassembly). Very few approaches to this problem have been proposed in the literature, and none is completely satisfactory. In this paper we present the hardware desig ...

Keywords: TCP/IP, asynchronous transfer mode, firewall, packet filtering, switch architecture

10 HTTP Cookies: Standards, privacy, and politics
David M. Kristol
November 2001 **ACM Transactions on Internet Technology (TOIT)**, Volume 1 Issue 2

Full text available: pdf(390.38 KB) Additional Information: full citation, abstract, references, citings, index terms

How did we get from a world where cookies were something you ate and where "nontechies"

h c g e cf c

were unaware of "Netscape cookies" to a world where cookies are a hot-button privacy issue for many computer users? This article describes how HTTP "cookies" work and how Netscape's original specification evolved into an IETF Proposed Standard. I also offer a personal perspective on how what began as a straightforward technical specification turned into a political flashpoint when it tried to address nontechn ...

Keywords: Cookies, HTTP, World Wide Web, privacy, state management

11 A case for two-way skewed-associative caches
André Seznec
May 1993 **ACM SIGARCH Computer Architecture News , Proceedings of the 20th annual international symposium on Computer architecture**, Volume 21 Issue 2
Full text available: pdf(975.20 KB) Additional Information: full citation, references, citings, index terms

12 A coherent distributed file cache with directory write-behind
Timothy Mann, Andrew Birrell, Andy Hisgen, Charles Jerian, Garret Swart
May 1994 **ACM Transactions on Computer Systems (TOCS)**, Volume 12 Issue 2
Full text available: pdf(3.21 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Extensive caching is a key feature of the Echo distributed file system. Echo client machines maintain coherent caches of file and directory data and properties, with write-behind (delayed write-back) of all cached information. Echo specifies ordering constraints on this write-behind, enabling applications to store and maintain consistent data structures in the file system even when crashes or network faults prevent some writes from being completed. In this paper we describe ...

Keywords: coherence, file caching, write-behind

13 Reducing cache misses using hardware and software page placement
Timothy Sherwood, Brad Calder, Joel Emer
May 1999 **Proceedings of the 13th international conference on Supercomputing**
Full text available: pdf(1.50 MB) Additional Information: full citation, references, citings, index terms

14 Query evaluation techniques for large databases
Goetz Graefe
June 1993 **ACM Computing Surveys (CSUR)**, Volume 25 Issue 2
Full text available: pdf(9.37 MB) Additional Information: full citation, abstract, references, citings, index terms, review

Database management systems will continue to manage large data volumes. Thus, efficient algorithms for accessing and manipulating large sets and sequences will be required to provide acceptable performance. The advent of object-oriented and extensible database systems will not solve this problem. On the contrary, modern data models exacerbate the problem: In order to manipulate large sets of complex objects as efficiently as today's database systems manipulate simple records, query-processi ...

Keywords: complex query evaluation plans, dynamic query evaluation plans, extensible database systems, iterators, object-oriented database systems, operator model of parallelization, parallel algorithms, relational database systems, set-matching algorithms, sort-hash duality

15 Let caches decay: reducing leakage energy via exploitation of cache generational behavior

Zhigang Hu, Stefanos Kaxiras, Margaret Martonosi

May 2002 **ACM Transactions on Computer Systems (TOCS)**, Volume 20 Issue 2

Full text available: pdf(873.03 KB) Additional Information: full citation, abstract, references, index terms

Power dissipation is increasingly important in CPUs ranging from those intended for mobile use, all the way up to high-performance processors for highend servers. Although the bulk of the power dissipated is dynamic switching power, leakage power is also beginning to be a concern. Chipmakers expect that in future chip generations, leakage's proportion of total chip power will increase significantly. This article examines methods for reducing leakage power within the cache memories of the CPU. Be ...

Keywords: Cache memories, cache decay, generational behavior, leakage power

16 Cache Memories

Alan Jay Smith

September 1982 **ACM Computing Surveys (CSUR)**, Volume 14 Issue 3

Full text available: pdf(4.61 MB) Additional Information: full citation, references, citings, index terms

17 Efficient simulation of caches under optimal replacement with applications to miss characterization

Rabin A. Sugumar, Santosh G. Abraham

June 1993 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 1993 ACM SIGMETRICS conference on Measurement and modeling of computer systems**, Volume 21 Issue 1

Full text available: pdf(1.26 MB) Additional Information: full citation, references, citings, index terms

18 Parallel execution of prolog programs: a survey

Gopal Gupta, Enrico Pontelli, Khayri A.M. Ali, Mats Carlsson, Manuel V. Hermenegildo

July 2001 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 23 Issue 4

Full text available: pdf(1.95 MB) Additional Information: full citation, abstract, references, citings, index terms

Since the early days of logic programming, researchers in the field realized the potential for exploitation of parallelism present in the execution of logic programs. Their high-level nature, the presence of nondeterminism, and their referential transparency, among other characteristics, make logic programs interesting candidates for obtaining speedups through parallel execution. At the same time, the fact that the typical applications of logic programming frequently involve irregular computatio ...

Keywords: Automatic parallelization, constraint programming, logic programming, parallelism, prolog

19 Tradeoffs in two-level on-chip caching

N. P. Jouppi, S. J. E. Wilton

April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST**

annual international symposium on Computer architecture, Volume 22 Issue 2

Full text available: pdf(1.16 MB) Additional Information: full citation, abstract, references, citings, index terms

The performance of two-level on-chip caching is investigated for a range of technology and architecture assumptions. The area and access time of each level of cache is modeled in detail. The results indicate that for most workloads, two-level cache configurations (with a set-associative second level) perform marginally better than single-level cache configurations that require the same chip area once the first-level cache sizes are 64KB or larger. Two-level configurations become even more import ...

20 Incremental computation via function caching

W. Pugh, T. Teitelbaum

January 1989 **Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on Principles of programming languages**

Full text available: pdf(1.55 MB) Additional Information: full citation, references, citings, index terms

h c g e cf c